

Mid-Range Computing Working Group Report

Defining the Future of Scientific Computing Resources at Berkeley Lab

CSAC and ITSD are working in partnership to determine the value of a Lab-wide Scientific Computing Resource for the future of LBNL scientific research

- Members
- Activities
- Schedule

The Working Group

CSAC

Paul Adams	<i>PBS</i>
Ali Belkacem	<i>CSD</i>
Alessandra Ciocio	<i>Physics (chair)</i>
Ken Downing	<i>LSD</i>
George Moridis	<i>ESD</i>
Doug Olson	<i>NSD</i>
John Staples	<i>AFRD</i>
Shaheen Tonse	<i>EETD</i>
Michel Van Hove	<i>MSD</i>
Tammy Welcome	<i>NERSC</i>

ITSD

Jon Bashor
Gary Jung
Jim Leighton
Yeen Mankin
Sandy Merola
Erik Richman
Dan Hawkes

Activities

- **Lecture Series Plan**

To provide information about the contribution of scientific computing to research

- **Web-based survey**

To determine the interest and needs of LBNL scientists in the area of MRC

- **Survey of MRC capability of other Labs**

To uncover success and failures of different MRC models

- **Cost estimate for different scenarios**

- **Study of current usage of computing resources at the Lab**

To find out where MRC would fit in to the current range of LBNL computers

- **Financial Model**

- **Formal Document**

To unify all documents and information

- **Publicity**

To be launched at the time of the lecture series

- **Schedule**

Lectures Series

To emphasize the use of large-scale computing in furthering scientific computations
To raise awareness of mid-range computing among LBNL scientists

Starting **September 12** through **October 3** (noon, 50-Auditorium, 90 minutes)

Each lecture:

- Will feature two scientific speakers, each talking 25-30 minutes on how high-end computing advanced their scientific programs
- Between the two speakers there would be shorter presentations
 - CSAC charter and MRC (Ali)
 - Berkeley Lab's New Cluster (Tammy)
 - Introducing the MRC survey form (Alessandra)

First lecture: ITSD view (Sandy) followed by Horst Simon or Bill McCurdy on the importance of high-end computing and vision of the future of scientific computing

WEB-based survey

To determine the interest and need of LBNL scientists in the area of MRC and identify key users

Erik Richman with Paul Adams and Shaheen Tonse produced a WEB-based questionnaire

URL's:

<http://www.lbl.gov/ITSD/CSAC/MRC> --- The main entry page

http://www.lbl.gov/ITSD/CSAC/MRC/CSAC_questionnaire.html --- The survey

We will also conduct additional one on one interviews with key users

All CSAC members are invited to test the questionnaire and give feedback

Survey of MRC capability of other Labs

Summary

Only ANL, LLNL, ORNL have some appearance of mid-range computing

Only LLNL actually has mid-range computing as a lab-wide resource

Scientists at many places have access to MRC that comes through some program but not as a lab-wide supported resource

Uniformly, past efforts to get users to pay for central computing has failed

LLNL has mechanisms for user programs to contribute funds to the central computing facility

Cost estimate for different scenarios

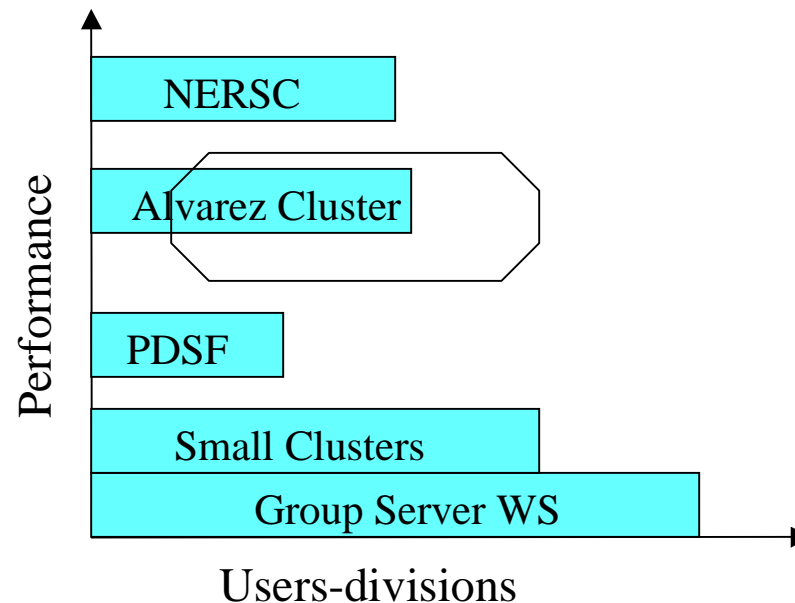
	Option 1 - Alvarez			Option 1a - Alvarez+			Option 2 - New Cluster			Option 3 - SMP		
	Year 1	Year 2	Year 3	Year 1	Year 2	Year 3	Year 1	Year 2	Year 3	Year1	Year2	Year3
MRC Purchase (no overhead)	0			500			700			1200		
plus procurement burden 3.9%	0			19.5			27.3			46.8		
plus materials handling 4.2%	0			21			29.4			50.4		
Procurement Total	0			540.5			756.7			1297		
Vendor support												
HW Maintenance (8x5) and SW Maintenance (15% of purchase cost)	0	52.5	105	0	90	180	35	35	35	180	180	180
Vendor Support Hotline	24	24	24	24	24	24	24	24	24	24	24	24
Additional software												
3rd party tools and applications												
Permanent license	251			251			251			129		
Annual software maintenance	48	48	48	48	48	48	48	48	48	24	24	24
Subtotal Vendor support and software	323	124.5	177	323	162	252	358	107	107	357	228	228
plus procurement burden 3.9%	12.60	4.86	6.90	12.60	6.32	9.83	13.96	4.17	4.17	13.92	8.89	8.89
Hardware and Software Support Total	336	129	184	336	168	262	372	111	111	371	237	237
Procurement Team Effort												
6 staff 0.5 FTE for 6 months (procurement, technical, benchmarks, etc...)				75			225			225		
Facilities Costs												
Base installation costs (seismic design, bracing, wiring)				20			60			60		
Power Distribution Unit including installation							60			60		
UPS including installation							60			60		
Space/Electricity	3.6	3.6	3.6	3.6	3.6	3.6	3.6	3.6	3.6	3.6	3.6	3.6
Staff Support (incl payroll and org burden)												
System Administration, Project Management, and Operations 3 FTE	450	486	525	450	486	525	450	486	525	450	486	525
User Services 1 FTE	150	162	175	150	162	175	150	162	175	150	162	175
Extra Cost Options List												
Platform Computing LSF (\$175/cpu purchase + \$50/cpu maintenance)	36	8	8	55.8	12.4	12.4	36	8	8	7.2	1.6	1.6
Applications Assistance 1 FTE	150	162	175	150	162	175	150	162	175	150	162	175
Total incl purchase, vendor support, software, staff and facilities	1125	951	1071	1705	994	1153	2323	933	998	2834	1052	1117
3 yr. Total Cost of Ownership	3147			3853			4254			5003		
25% planning margin	3933			4816			5317			6254		

Study of current usage of computing resources at the Lab

To find out where an MRC would fit in to the current range of LBNL computers from Group Server Workstations to NERSC

By estimating the scale of computation at the Lab and comparing the relative power of the various computing platforms

MRC may be the right replacement for the ending T3E program



Financial Model

- Users will probably not want a recharge funding model for a joint MRC machine at LBNL
- An alternative is for central ownership - users pay for CPU time by contributing hardware
- The Laboratory could provide add-ons such as software

A mixed funding model may be:

60% of the machine bought by divisions
30% by overhead and
10% provided by special LBNL funding

- The institutional support may consist of 2-3 FTEs (maintenance and programming help)
- Divisions that do not use the system would still contribute through the 30% overhead charge
- Some divisions may be elected to participate at a later date
- Operating system software would come out of overhead funds

If this machine had a similar configuration as NERSC, it would serve as a stepping-stone to NERSC

Formal Document

A 10-page document containing an introduction and including in several sections the information we have uncovered. (Jon Bashor and Dan Hawkes)

Sections:

- Cost estimate
- High Performance Computing Facility use by LBNL users
- Draft of survey that will be presented at the lecture series
- Plan for the lecture series
- White paper (under revision)
- Survey of key users of scientific computing at LBNL
- Information on MRC at other labs (success and failures)

In progress:

- A financial strategy for funding an MRC facility
probably looser than the 60/30/10%
- A chart showing computer capability vs. number of {users, groups, users*GFlops}

Schedule

May

Produce formal document

June

Presentation at Division Directors Retreat

Sept/Oct

Launch Publicity Campaign

Lecture series

Web survey release towards the end of lecture series

Nov/Dec

Analyze survey results

Interview with key users

One-day retreat to define system

January

Present at CSAC the results and finalize recommendation